Performance of LDPC Decoders with Missing Connections

Linjia Chang, *Student Member, IEEE,* Avhishek Chatterjee, and Lav R. Varshney, *Senior Member, IEEE*

arXiv:1605.00317v4 [cs.IT] 9 Jan 2017

***Abstract*—Due to process variation in nanoscale manufacturing, there may be permanently missing connections in information processing hardware. Due to timing errors in circuits, there may be missed messages in intra-chip communications, equivalent to transiently missing connections. In this work, we investigate the performance of message-passing LDPC decoders in the presence of missing connections. We prove concentration and convergence theorems that validate the use of density evolution performance analysis. Arbitrarily small error probability is not possible with missing connections, but we find suitably defined decoding thresholds for communication systems with binary erasure channels under peeling decoding, as well as binary symmetric channels under Gallager A and B decoding. We see that decoding is robust to missing wires, as decoding thresholds degrade smoothly. Moreover, there is a stochastic facilitation (SF) effect in Gallager B decoders with missing connections. We also conduct finite-length simulations, compare the decoding sensitivity to channel noise and to missing wiring, and perform preliminary error-tolerant manufacturing yield analysis.**

***Index Terms*—Decoding, error analysis, message passing, wiring, stochastic facilitation**

I. Introduction

LOW-DENSITY parity-check (LDPC) codes are prevalent due to their performance near the Shannon limit with message-passing decoders that have efficient implementation [2]. With the end of CMOS scaling nearing, there is interest in nanoscale circuit implementations of decoders, but this introduces concerns that process variation in manufacturing may lead to interconnect patterns different than designed [3]–[5], especially under self-assembly [6], [7]. Yield on manufactured chips deemed perfectly operational is small—reports indicate 1–15% of circuit elements such as wires, switches, and transistors are defective [7]—leading to rather expensive industrial waste [8]. Changing the paradigm of circuit functionality from perfection to some small probability α of missing wires may eliminate much wastage and so it is of interest to characterize chips with permanently missing connections to determine suitable error tolerances.

Process variation in manufacturing also causes fluctuation in device geometries, which might prevent them from meeting timing constraints [9], especially in future nanoscale technologies like carbon nanotube circuits where device geometry control is especially difficult. Such timing errors lead to missed messages in intra-chip communications, equivalent to transiently missing connections. Connections can also be missing transiently in programmable LDPC decoders [10]. It is thus also of interest to characterize decoders with transiently missing connections.

However, most fault-tolerant computing research assumes the circuit is constructed correctly and is concerned only with faults in computational elements. Peter Elias noted the following [11], but it remains true today:

> J. Von Neumann has analyzed computers whose unreliable elements are majority organs—crude models of a neuron. Shannon and Moore have analyzed combinational circuits whose components are unreliable relays. Both papers assume that the wiring diagram is correctly drawn and correctly followed in construction, but that computation proper is performed only by unreliable elements.

Such assumptions of fault-free circuit construction need to be reevaluated and performance analysis of computation with such wiring faults needs to be carried out. The only work we are aware of in fault-tolerant computing theory that briefly discusses wiring errors is [12, Ch. 9.2].

We had previously extended the method of density evolution to decoders with faults in the computational elements and showed it is possible to communicate with arbitrarily small error probability with noisy Gaussian belief propagation [13]. Asymptotic characterizations were also determined for Gallager A [13] and Gallager B decoders with transient noise [14]–[16], energy optimization [17], and both permanent and transient noise [18]. Noisy decoding [19]–[25], and general noisy belief propagation, not necessarily in decoding [26], [27], have also been studied. Recent studies show bit-flipping decoders with data-dependent gate failures can achieve zero error probability [25], [28], but with a subset of computation hardware that is reliable and no wiring diagram errors.

Rather than noise in computational elements, here we analyze the performance of message-passing decoders with missing connections and show that appropriately defined decoding thresholds are robust, in the sense of degrading smoothly. This is true for both transiently and permanently missing connections in message-passing decoding circuits. In certain settings, missing connections actually improve performance, resulting in stochastic facilitation (SF).[1]

This work was presented in part at the 2016 IEEE International Symposium on Information Theory [1].

This work was supported in part by Systems on Nanoscale Information fabriCs (SONIC), one of the six SRC STARnet Centers, sponsored by MARCO and DARPA.

L. Chang, A. Chatterjee, and L. R. Varshney are with the Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, Urbana, IL 61801 USA (e-mail: {lchang10, avhishek, varshney}@illinois.edu).

[1]SF in decoding was observed with transient errors in computation, rather than with missing connections, initially in memory recall [27], [29] and then in communications [23], [30].

A key difference between noisy computational elements and missing connections is that circuit technology enables detection of missing connections, see Sec. II-C. This allows for simple adaptations of decoding algorithms, yielding better decoding performance under missing connections than under noisy components. A notable manifestation of this difference is in the so-called decoder useful region. For transient or permanent noise, there is a strictly positive lower bound for the useful region, below which the channel output is actually better than the decoded version since the internal decoder noise makes things worse. For missing connections, there is no such lower boundary since the decoder asymptotically never degrades performance from the raw channel error rate.

The celebrated results of Richardson and Urbanke [31] developed density evolution for analyzing message-passing decoders for LDPC codes that are correctly wired. Here we extend those results, so we can use the density evolution technique to characterize symbol error rate P_e, measuring the fraction of incorrectly decoded symbols at the end of message-passing decoding, even when the decoder has missing connections. We also show that the performance of decoders with transiently and with permanently missing connections are asymptotically equivalent. Traditionally [31], there are thresholds for channel noise level ε below which P_e can be driven to 0 with increasing blocklength n. Unfortunately with missing connections in message-passing decoders, P_e cannot be driven to 0 in general without a significant modification of the decoding algorithm. Thus, following [13], we let η upper bound the final error probability achievable by decoders with missing connections after many iterations ℓ and give thresholds to ε, below which $\lim_{\ell\to\infty} P_e^{(\ell)} \leq \eta$ under density evolution.

We perform sensitivity analysis of density evolution to give insight into whether manufacturing or operational resources are more important in communication infrastructures. We also comment on how our results inform semiconductor manufacturing yield analysis under the new paradigm of allowing some level of wiring error. To demonstrate the practical utility of density evolution analysis, we also perform finite-length simulations of decoders with missing connections.

Sec. II discusses models of codes, channels, and LDPC decoders with both transiently and permanently missing connections, with a particular focus on hardware modeling. Sec. III develops tools including concentration and convergence theorems that provide validity to density evolution analysis. Secs IV, V, and VI analyze the peeling decoder on the binary erasure channel (BEC) and the Gallager A and Gallager B decoders on binary symmetric channel (BSC) using density evolution, characterizing P_e with missing connections. Sec. VII connects our work to practice through sensitivity analysis, finite-length simulations, and manufacturing yield analysis. Sec. VIII concludes by pointing out directions for further investigation.

II. Background

In this section we describe message-passing LDPC decoders with missing connections. We define the code and channels considered in this work and construct fault-free and missing-wire decoder models for characterization later.

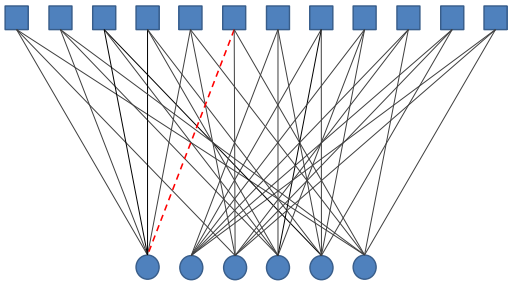

Fig. 1. Tanner graph of a $(3, 6)$ regular LDPC code, with a missing wire for a corresponding message-passing decoder highlighted with a dashed line.

A. Ensemble of LDPC Codes and Channel

We are concerned with the standard LDPC code ensemble G^n, both regular and irregular. First consider $(d_\mathrm{v}, d_\mathrm{c})$-regular LDPC codes of length n, which can be defined by a bipartite Tanner graph with n variable nodes of degree d_v in one set, and $nd_\mathrm{v}/d_\mathrm{c}$ check nodes of degree d_c in the other set (see Fig. 1). For irregular codes $G^n(\lambda, \rho)$, the degree distribution of variable and check nodes are denoted by functions $\lambda(x) = \sum_{d=2}^{\infty} \lambda_d x^{d-1}$ and $\rho(x) = \sum_{d=2}^{\infty} \rho_d x^{d-1}$, where λ_d and ρ_d specify the fraction of edges in the graph that are connected to nodes with degree d. The variable nodes hold the codeword messages, and the check nodes enforce the constraints among variable nodes according to the code design. We consider this binary linear code ensemble as defined over the alphabet $\{\pm 1\}$. Although this section is general, for convenience, let us think of the communication channel as either BSC with output alphabet $\{\pm 1\}$ or BEC with output alphabet $\{\pm 1, ?\}$.

B. Fault-Free Message-Passing Decoder

The decoder operates by passing messages iteratively over the edges in the Tanner graph of the code. The implementation of such message-passing decoders in hardware follows the construction of the same Tanner graph too. We define a variable-to-check node message $u_{v\to c}$ and a check-to-variable node message $u_{c\to v}$. Message $u_{v'\to c'}$ from variable node v' to check node c' is often computed based on all incoming $u_{c\to v'}$ messages, where $c \in N(v')$ is a neighboring node of v' and $c \neq c'$. For peeling, Gallager A, and Gallager B decoders, message $v_{c'\to v'}$ from check node c' to variable node v' is the product of all incoming $u_{v\to c'}$ messages, where $v \in N(c')$ is a neighboring node of c' and $v \neq v'$.

C. Missing Connections

As discussed in Sec. I, there are two types of missing connections: permanent missing connections caused by breaks in interconnects and transient missing connections caused mainly by timing errors in intra-chip communication due to geometry variation in circuitry. Although specific statistical characterization is not reported in the semiconductor industry, process variation in manufacturing leads to both kinds of errors and can be fairly prevalent [7].

For a given decoder circuit, permanent failure is modeled by removing each connection between variable and check nodes

with probability α independently from others, before decoding starts. These connections are never active once removed. Conversations with circuit designers suggest that when an interconnect is broken in LDPC decoders implemented in a variety of device technologies, the measured signal voltage at this open-ended wire is neither low (0–$0.3V_{dd}$) nor high ($0.7V_{dd}$–V_{dd}); it is an intermediate floating value varying in the range ($0.3V_{dd}$–$0.7V_{dd}$) that can be differentiated from high/low values. Hence we model it as an erasure symbol, "?".

For the transiently missing connection setting, each connection between variable and check nodes is removed independently from others with probability α at each decoding iteration. Transiently missing connections may occur due to timing error from incorrect geometry: consider a misalignment of synchronization when one branch of signal arrives after the computation at the destination node has started, especially among those circuit implementations that do not store the last signal sample. Transiently missing connections might similarly happen in programmable LDPC decoder architectures [10]. Due to the difficulty in controlling device geometries, future carbon nanotube circuits are projected to have a significant number of these transient missing connections. Again we model as an erasure symbol, "?".

For notational convenience, let us restrict attention to decoders with messages in $\{\pm 1, ?\}$, but again concentration and convergence results demonstrated in Sec. III are general. Motivated by different concerns, [32, Ex. 4.86] considered erasures in decoder messages as a representation of confidence, whereas [33] considered erasures as a way to capture check node or variable node failures in belief propagation.

III. Performance Analysis Tools

We now present mathematical tools to simplify the performance analysis of LDPC decoders with missing connections. In particular, we establish symmetry conditions for binary codes for easy analysis, and concentration and convergence results that endow the density evolution method with significance. Such results can be applied not only to decoders with binary messages, but also with larger message sets.

A. Restriction to All-One Codeword

Under certain symmetry conditions of the code, the communication channel, and the message-passing decoder, the probability of error is independent of the transmitted codeword.

C1. **Code Symmetry:** Code is a binary linear code.
C2. **Channel Symmetry:** Channel is a binary memoryless symmetric channel [32, Def. 4.3 and 4.8].
C3. **Check Node Symmetry:** If incoming messages of a check node are multiplied by $\{b_i \in \{\pm 1\}\}$, then the computed message is multiplied by $\prod_i b_i$.
C4. **Variable Node Symmetry:** If the sign of each incoming message is flipped, the sign of the computed message is also flipped.

Proposition 1: Under conditions C1–C4, in the presence of transiently or permanently missing connections, the probability of error of a message passing decoder is independent of the transmitted codeword.

Proof: First consider mapping the erasure message "?", sent when a connection is missing, to 0; thus the check-to-variable and variable-to-check messages are the messages computed at check node and variable node, respectively, multiplied by either 1 (connection exists) or 0 (missing connections). It follows that messages passed between check and variable nodes satisfy the respective symmetry conditions [32, Def. 4.82]. Hence, the result follows from [32, Lem. 4.92]. ■

In the sequel, we restrict the analysis of all models to the all-one codeword.

B. Concentration around Ensemble Average

We now show that the performance of LDPC codes decoded with missing-connection decoders stays close to the expected performance of the code ensemble for both transiently and permanently missing wires. The approach follows [31] and is based on constructing an exposure Martingale, obtaining bounded difference constants, and using Azuma's inequality.

Fix the number of decoding iterations at some finite ℓ and let Z be the number of incorrect values held among all $d_{\mathrm{v}}n$ edges at the end of the ℓth iteration for a specific choice of code, channel noise, and decoder with missing wires. Let $E[Z]$ be the expectation of Z. Thm. 1 holds for decoders with both transiently and permanently missing connections.

Theorem 1 (Concentration around Expected Value): There exists a positive constant $\beta = \beta(d_{\mathrm{v}}, d_{\mathrm{c}}, \ell)$ such that for any $\varepsilon > 0$,

$$\Pr[|Z - E[Z]| > n d_{\mathrm{v}} \varepsilon/2] \leq 2e^{-\beta \varepsilon^2 n}.$$

Proof: See Appendixes B and C for permanent and transient missing connections, respectively. ■

Recall Doob's Martingale construction from [31], and the bounded difference constants for exposing channel noise realizations and the realized code connections, together with Azuma's inequality. The main difference between the Martingale construction here and [31] is in the bounded differences due to the additional randomness from missing connections.

For permanently missing connections, one can think of the final connection graph being sampled from an ensemble of irregular random graphs with binomial degree distribution with average degrees $(1-\alpha)d_{\mathrm{c}}$ and $(1-\alpha)d_{\mathrm{v}}$, bounded by maximum degrees d_{c} and d_{v}. Hence, the result follows from the result for correctly-wired irregular codes [31].

For transiently missing connections, the Martingale is constructed differently. Here instead of edges, for ℓ iterations, we sequentially expose the realization of edges at different iterations. Similar to [13] for transient noise in computational elements, the Martingale difference is bounded using the maximum number of edges over which a message can propagate in ℓ iterations, by unwrapping a computation tree.

Note that β is smaller for transient than permanent miswiring. The theorem extends directly to irregular LDPC codes.

C. Cycle-Free Case

We now show that the average performance of an LDPC code ensemble converges to an associated cycle-free tree structure, unwrapping a computation tree as in [31].

For an edge whose connected neighborhood with depth 2ℓ is cycle-free, let q denote the expected number of incorrect values held along this edge at the end of ℓth decoding iteration. The expectation is taken over the choice of code, the messages received from the channel, and the realization of the decoder with missing wires. The theorems hold for both transiently and permanently missing connections.

Theorem 2 (Convergence to Cycle-Free Case): There exists a positive constant $\gamma = \gamma(d_{\mathtt{v}}, d_{\mathtt{c}}, \ell)$ such that for any $\varepsilon > 0$ and $n > \frac{2\gamma}{\varepsilon}$,

$$|E[Z] - nd_{\mathtt{v}}q| < nd_{\mathtt{v}}\varepsilon/2.$$

Proof: The proof is identical to [31, Thm. 2], since introducing missing connections in a cycle-free tree structure does not change its cycle-free property. ∎

Theorem 3 (Concentration around Cycle-Free Case): There exist positive constants $\beta = \beta(d_{\mathtt{v}}, d_{\mathtt{c}}, \ell)$ and $\gamma = \gamma(d_{\mathtt{v}}, d_{\mathtt{c}}, \ell)$ such that for any $\varepsilon > 0$ and $n > \frac{2\gamma}{\varepsilon}$,

$$\Pr[|Z - nd_{\mathtt{v}}q| > nd_{\mathtt{v}}\varepsilon] \leq 2e^{-\beta n\varepsilon^2}.$$

Proof: Follows directly from Thms. 1 and 2. ∎

This concentration result holds for all message-passing decoders with missing connections. In the sequel, we consider special cases of peeling, Gallager A, and Gallager B decoders.

D. Density Evolution

With the concentration around the cycle-free case, it is clear that the symbol error rate P_e of message-passing decoders with missing connections can be characterized with the density evolution technique. Let $P_e^{(\ell)}(g, \varepsilon, \alpha)$ be the error probability of decoding a code $g \in G^n$, after the ℓth iteration, where ε is the channel noise parameter, and α is the decoder missing wire probability. Density evolution evaluates the term:

$$\bar{P}_e^{\,(\ell)} = \lim_{n\to\infty} E[P_e^{(\ell)}(g, \varepsilon, \alpha)].$$

The expectation is over the choice of code, channel noise realization, and missing wire realization.

Based on the proof of Thm. 2, we claim that the decoding error probability at any iteration ℓ for transiently and permanently missing connections, $\bar{P}_{eT}^{(\ell)}$ and $\bar{P}_{eP}^{(\ell)}$, become identical with the increase of the girth as blocklength n increases. In particular, in density evolution the state variable $x_{\ell+1}$ is computed based on the x_ℓ values of nodes immediately below in the infinite tree. Each connection in the tree is encountered only once. In case of permanent failures each connection is present in the decoder with probability $1 - \alpha$, whereas for transient failures each connection is present at any iteration with probability $1 - \alpha$. But, for a given code symbol, its intrinsic messages traverse a particular edge only once if the LDPC graph is a tree. Thus the messages experience the same statistical effect under permanent and transient failures. This results in the same probability of error under both failures.

Theorem 4: For any arbitrarily small $\delta = \delta(d_{\mathtt{v}}, d_{\mathtt{c}}, \ell) > 0$, $\sigma > 0$, and $\ell \geq 0$:

$$\Pr[|\bar{P}_{eT}^{(\ell)} - \bar{P}_{eP}^{(\ell)}| \geq \sigma] \leq \delta.$$

Proof: First, let $N_{\vec{e}}^{2\ell}$ be the neighborhood of an edge $\vec{e}$ with depth 2ℓ in the decoding graph. Define the event A_N as $N_{\vec{e}}^{2\ell}$ is not tree-like. It is shown that for a positive constant $\tau = \tau(d_{\mathtt{v}}, d_{\mathtt{c}}, \ell)$, $\Pr[A_N] \leq \frac{\tau}{n}$ [31, Thm. 2]. It implies the probability of exposing an edge multiple times decreases with increasing blocklength n at any iteration ℓ. Following the edge exposing procedure, $\bar{P}_{eT}^{(\ell)}$ and $\bar{P}_{eP}^{(\ell)}$ differ only when any edge $\vec{e}$ is exposed multiple times and the presence of $\vec{e}$ in the two decoding graphs with permanently and transiently missing connections differs. Hence, $\Pr[|\bar{P}_{eT}^{(\ell)} - \bar{P}_{eP}^{(\ell)}| \geq \sigma] = \Pr[|\bar{P}_{eT}^{(\ell)} - \bar{P}_{eP}^{(\ell)}| \geq \sigma \big| A_N] \Pr[A_N] + \Pr[|\bar{P}_{eT}^{(\ell)} - \bar{P}_{eP}^{(\ell)}| \geq \sigma \big| {A_N}^c] \Pr[{A_N}^c]$. Since $\Pr[|\bar{P}_{eT}^{(\ell)} - \bar{P}_{eP}^{(\ell)}| \geq \sigma \big| {A_N}^c] = 0$, we can show $\Pr[|\bar{P}_{eT}^{(\ell)} - \bar{P}_{eP}^{(\ell)}| \geq \sigma] \leq 1 \cdot \Pr[A_N] \leq \frac{\tau}{n}$. As $n \to \infty$, this probability $\frac{\tau}{n} = \delta$ approaches 0. ∎

In the sequel, no distinction is made between the analysis for transiently and permanently missing connection cases.

E. Decoder Useful Region and Thresholds

Usually density evolution converges to a certain stable fixed point with increasing number of iterations ℓ. We define this fixed point as:

$$\bar{P}_e^{\,(\infty)} = \lim_{\ell\to\infty} \bar{P}_e^{\,(\ell)} = \lim_{\ell\to\infty}\lim_{n\to\infty} E[P_e^{(\ell)}(g, \varepsilon, \alpha)].$$

In order to decide when to use a decoder with missing connections, a *useful decoder* is defined. A decoder is said to be useful and should be used instead of taking the codeword directly from the channel without decoding, if the asymptotic decoding error probability satisfies [13]:

$$\bar{P}_e^{\,(\infty)} < \bar{P}_e^{\,(0)} = \varepsilon.$$

The useful region of a decoder is defined as the set of parameters, in our case (ε, α), that satisfies the above condition. Note that in case of transient computation noise where computation is erroneous with probability $\tilde{\alpha}$ [13], there are $(\varepsilon, \tilde{\alpha})$ such that $\bar{P}_e^{\,(\infty)} > \varepsilon$. But under missing connections, for peeling, Gallager A, and Gallager B decoders, $\bar{P}_e^{\,(\infty)} \leq \varepsilon$ for any (ε, α). This is because these decoders do not propagate erroneous messages under missing connections and hence cannot degrade symbols received from the channel. When decoding with a fault-free decoder where $\alpha = 0$, there exists an ε^* below which the final decoding error probability goes to 0 and a much larger value otherwise. We will see in the following sections that $\bar{P}_e^{\,(\infty)}$ does not go to zero for positive α, but a threshold phenomenon still exists.[2] For every fixed α, there exists a channel noise decoding threshold ε^*, below which the final error probability $\bar{P}_e^{\,(\infty)}$ goes to a small value η. We call decoders that can achieve $\bar{P}_e^{\,(\infty)}$ that is lower than this small value η-reliable, and the channel noise level beyond which the decoder is η-reliable, the decoding threshold ε^* [13]:

$$\varepsilon^*(\eta, \alpha) = \sup\big\{\varepsilon \in [0, 0.5] \big| \bar{P}_e^{\,(\infty)} \text{ exists and } \bar{P}_e^{\,(\infty)} < \eta\big\}.$$

[2] In general for $\bar{P}_e^{\,(\infty)}$ to go to 0 in a faulty decoder, one needs to substantially change the decoder or to have a structural relationship between data and errors [17], [25], [28].

IV. Peeling Decoder

Consider the peeling decoder for communication over a BEC with alphabet $\{\pm 1, ?\}$. The check node computation is a product of all messages ± 1 it receives from neighboring variable nodes if none is "?", otherwise an erasure symbol "?" is sent. The variable node computation is to send any ± 1 symbol received either from the other check nodes or from the channel, otherwise send "?". When the connection between two nodes is missing, the message exchanged is equivalent to "?", so peeling extends naturally to decoders with missing connections. Note that this decoder satisfies the symmetry conditions C1–C4, so we can use density evolution assuming the all-one codeword was transmitted.

Although high-level intuition would suggest that the performance of decoding would degrade for any code and any decoder with missing connections, this is not the case as we later show for the Gallager B decoder. For the peeling decoder, the intuition holds and can be formalized using coupling techniques and the fact that peeling decoders never propagate erroneous messages.

Lemma 1: For any LDPC code g with an arbitrary but finite blocklength, after a finite number of decoding iterations ℓ, for both permanently and transiently missing connections, the symbol error probability $P_e^{(\ell)}(g, \varepsilon, \alpha)$ increases monotonically with α for a given ε.

Proof: The proof for monotonicity of $P_e^{(\ell)}(g, \varepsilon, \alpha)$ follows by simple coupling arguments. For a specific LDPC code, consider two different missing connection probabilities α_1 and α_2, where $\alpha_1 < \alpha_2$. Then, we couple the two missing connection processes as follows. Remove the wires with probability α_1, and from this check-variable connection graph, remove each of the remaining connections with probability $\alpha_2 - \alpha_1$. This gives a second missing connection process. It can be checked that the probability of missing connection in the second process is α_2. Thus we can couple the missing connection processes to get a sample path dominance of connections. In this coupling, any realization of α_2 process has more missing connections than that of α_1.

Now consider the probability of correctly decoding any bit i. Note that in peeling decoders, no erroneous messages are exchanged between check and variable nodes; only correct messages and erasures are passed along wires. A variable node v_i holding message bit i can be decoded correctly if either the received bit is correct, or the received bit is an erasure but v_i receives a correct message through a path on the computation tree passing through one of its check nodes. The probability the received bit is correct is the same in case of both α_1 and α_2. So, let us compare the other probability. Now, by coupling as any realization of α_1 has more connections than α_2, if a correct message reaches i following a path in the α_2 graph, then that path also exists in the α_1 graph. Thus, the event of receiving a correct message in case of α_2 is a subset of that of α_1. This proves monotonicity of correct decoding probability and so missing connections only degrade performance. ∎

A similar coupling argument yields an ordering relationship with respect to channel erasure probability ε for a given α.

A. Density Evolution Equation

First, recall that the peeling decoding algorithm allows $\{\pm 1, ?\}$ to be sent, where "?" stands for an erasure caused by either the channel noise or a missing connection. In this case, the decoder only outputs either the correct message or an erasure symbol.

Consider a regular $(d_\mathrm{v}, d_\mathrm{c})$ LDPC code, BEC channel with parameter ε, and each wire independently disconnected with probability α. Let $x_0, x_1, \ldots, x_\ell$ denote the fraction of erasures existing in the code at each decoding iteration. The original received message from the channel is erased with probability ε, so $P_e^{(0)}(\varepsilon, \alpha) = x_0 = \varepsilon$.

Let q_{in} be the probability that a node receives an erasure, and q_{out} be the probability that a node sends out an erasure. At a variable node, the probability that a given internal incident variable will be erased is the probability that both the external incident variable is erased and all other $d_\mathrm{v} - 1$ nodes are either disconnected or connected but erased.

$$\begin{aligned} q_{out} &= x_0 \sum_{i=0}^{d_\mathrm{v}-1} \binom{d_\mathrm{v}-1}{i} \alpha^i [q_{in}(1-\alpha)]^{(d_\mathrm{v}-1)-i} \\ &= \varepsilon[\alpha + (1-\alpha)q_{in}]^{d_\mathrm{v}-1}. \end{aligned}$$

At a check node, the probability that a given incident variable will not be erased is the probability that all $d_\mathrm{c} - 1$ other internal incident variables are not erased or disconnected. So the probability that a message is erased is

$$q_{out} = 1 - [(1-q_{in})(1-\alpha)]^{d_\mathrm{c}-1}.$$

Hence, the density evolution of the fraction of erasure between two consecutive decoding iterations is

$$x_{\ell+1} = \varepsilon\big[\alpha + (1-\alpha)\big(1 - [(1-x_\ell)(1-\alpha)]^{d_\mathrm{c}-1}\big)\big]^{d_\mathrm{v}-1}.$$

The density evolution result can be extended to irregular LDPC codes:

$$x_{\ell+1} = \varepsilon\lambda\Big(\alpha + (1-\alpha)\big(1 - \rho[(1-x_\ell)(1-\alpha)]\big)\Big).$$

Let $f_{DE}(x_\ell, \varepsilon, \alpha) = x_{\ell+1}$ be the recursive update function for the fraction of erasure, where $0 \le \varepsilon < 0.5$ and $0 \le \alpha \le 1$ is the domain of interest.

B. Fixed Points

The density evolution function f_{DE} is non-decreasing in each of its arguments, given the other two. Thus, a monotonicity result similar to [32, Lem. 3.54] holds. This also implies a convergence result for x_ℓ, similar to [32, Lem. 3.56]. So, for a given α and ε, x_ℓ converges to the nearest fixed point of $x = f_{DE}(\varepsilon, x, \alpha)$. Due to this existence of the fixed point, we can characterize the error probability when the decoding process is finished. The fixed points can be found by solving for the real solutions to the polynomial equation

$$x - \varepsilon\lambda\Big(\alpha + (1-\alpha)\big(1 - \rho[(1-x)(1-\alpha)]\big)\Big) = 0. \quad (1)$$

We now prove that the decoding error probability is strictly positive by showing that $x = 0$ is not a fixed point in (1).

Lemma 2: For any irregular code ensemble $C^{\infty}(\lambda, \rho)$, there exists a $\delta > 0$, such that the probability of error $P_e^{(\infty)}$ satisfies $P_e^{(\infty)} > \varepsilon\lambda(1-(1-\alpha)\rho(1-\alpha)) > \delta > 0$.

Proof: Since x_ℓ is monotonic, if $x_0 \leq x_1$ then for any ℓ, $x_{\ell+1} \geq x_\ell \geq x_{\ell-1}$. Now, for $x_0 = 0$, by substituting this value in f_{DE},

$$x_1 = f_{DE}(0, \varepsilon, \alpha) = \varepsilon\lambda(1-(1-\alpha)\rho(1-\alpha)) > 0 = x_0.$$

This implies that $\lim_{\ell\to\infty} x_\ell \geq f_{DE}(0, \varepsilon, \alpha)$, for $x_0 = 0$. But, as x_ℓ converges to the fixed point nearest to x_0 in the direction of monotonicity, $x = 0$ is not a fixed point and there is no fixed point in $(0, f_{DE}(0, \varepsilon, \alpha))$ for any $\varepsilon, \alpha > 0$. Thus we have $P_e^{(\infty)} > 0$. ∎

Since this lemma shows all fixed points of the density evolution equation are greater than $\varepsilon\lambda(1-(1-\alpha)\rho(1-\alpha))$, decoding error probability cannot be taken to zero. But this does not mean that the decoder is not useful. In fact it is always better to use the decoder, even when there are missing connections, rather than just taking corrupted symbols from the channel directly, since the peeling decoder never has incorrect messages. We can see this using the monotonicity of $f_{DE}(x, \varepsilon, \alpha)$ in each of its arguments, given the other two. For any channel and code, $x_0 = \varepsilon$, and it follows from f_{DE} that $x_1 = f_{DE}(x_0, \varepsilon.\alpha) \leq \varepsilon$. Hence $x_\ell \leq \varepsilon$, for all ℓ, and $P_e^{(\infty)} \leq \varepsilon$. This is in sharp contrast to decoders with computation noise, where decoder output can be strictly worse than channel output [13].

C. Performance Analysis

In the previous section, we developed the recursive function to characterize the final error probability achieved by a peeling decoder with missing wires. Now we want to characterize the performance of such decoders.

For a peeling decoder, when $\varepsilon = 0$, the error probability stays at 0 regardless of the quality of the decoder. When $\alpha = 0$, it has been shown that there exists decoding threshold on the channel noise ε, below which the final error probability can be driven to 0 with the increase of decoding iterations [31]. For the following analysis, we consider the system when $\varepsilon > 0$ and $\alpha > 0$. Ideally, we want the error probability to be driven to 0, but as demonstrated in Lem. 2, this is impossible. Here we use the weaker notion of η-reliability defined in Sec. III-E, where η limits the final decoding error probability P_e.

Fig. 2 shows the final symbol error rate of decoding a $C^{\infty}(3, 6)$ LDPC code under peeling decoding with various missing connection probabilities α over BEC(ε). It can be seen that given α, there exists a threshold in channel noise level where a phase transition in P_e happens. Fig. 3 illustrates such thresholds with the change of α under different small η-reliable constraints. An interesting phenomenon to notice in the decoding threshold is that there also exists a phase transition with the change of the decoder missing connection probability α. With the increase of α, for a fixed η-reliable decoder with missing connections, the decoding threshold first decreases linearly, and then exhibits more rapid decrease before convergence to zero.

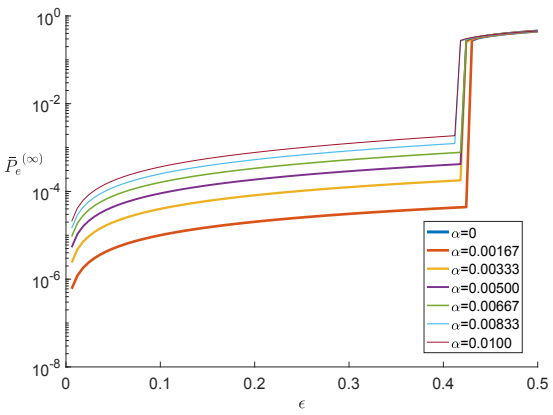

Fig. 2. Final symbol error rate of decoding a $C^{\infty}(3, 6)$ LDPC code under peeling decoding algorithm with various missing connection probability α over BEC.

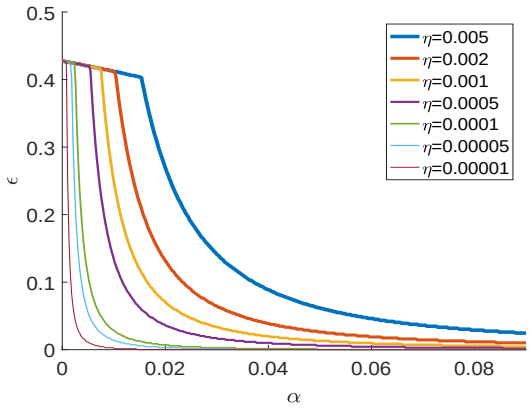

Fig. 3. Channel threshold of decoding a $C^{\infty}(3, 6)$ LDPC code under peeling decoding algorithm over BEC for different given final error η-thresholds.

V. Gallager A Decoder

Consider a fault-free Gallager A decoder for communication over a BSC. The messages are passed along the edges in the corresponding Tanner graph during decoding. A check node computes the product of incoming variable-to-check node messages $\{u_{v\to c}\}$; a variable node decides to flip the message from channel y_v if all of the incoming check-to-variable node messages are $-y_v$ [2].

With missing connections, the check node computation is not defined if an input is unknown ("?"). The product computed at the check node is the modulo-2 sum of all incoming messages to ensure that the parity constraints of the code are satisfied. When one of the bits involved in the parity is unknown, that parity check is no longer informative. This is because any bit of a linear code is equally likely to be ± 1 (as complementing a binary codeword gives a codeword). So, for decoders with missing connections we make a natural adaptation: $u_{c\to v}$ = "?" if any of the incoming messages is "?". We also make a natural adaptation for variable node computation: $-y_v$ is sent if more than one non-erasure check node messages are $-y_v$, and y_v is sent otherwise.

When it comes to Gallager A decoding over BSC, the messages being passed between nodes may carry erroneous

information, unlike the peeling decoder, where the messages are either correct or erasure. So, for a sample path realization of channel and missing connections, a missing connection may prevent propagation of erroneous messages. Hence, unlike the peeling decoder, it is not apparent that there exists a stochastic dominance result like Lem. 1 between two different probabilities of missing connections. As fault-free decoding with the Gallager A algorithm satisfies conditions C1–C4, we can restrict analysis to the all-one codeword.

A. Density Evolution Equation

We find the probability for a variable node to compute -1 at iteration $\ell+1$, in terms of x_ℓ. We consider a regular $(d_\mathsf{v}, d_\mathsf{c})$ LDPC code and the adaptation of Gallager A decoding with erasure symbols for missing connections.

First note that since a BSC only outputs ± 1, a variable node never computes "?" with the Gallager A adaptation, even though it may receive (due to connection failure or check-node computes "?") or send the erasure symbol "?" (only due to connection failure).

The probability that a check node computation is -1 is:

$$\begin{aligned}
&\Pr\big\{\text{all } (d_\mathsf{c}-1) \text{ variable nodes are connected and send odd number of } -1\big\}\\
&= (1-\alpha)^{d_\mathsf{c}-1} \Pr\{\text{odd number of } (d_\mathsf{c}-1) \text{ nodes send } -1\}\\
&= (1-\alpha)^{d_\mathsf{c}-1} \tfrac{(1-(1-2x_\ell)^{d_\mathsf{c}-1})}{2},
\end{aligned}$$

where the last line follows using results from [2, Sec. 4.3].

The probability that a check node computation is $+1$ is:

$$\begin{aligned}
&\Pr\big\{\text{all } (d_\mathsf{c}-1) \text{ variable nodes are connected and send even number of } -1\big\}\\
&= (1-\alpha)^{d_\mathsf{c}-1} \Pr\{\text{even number of } (d_\mathsf{c}-1) \text{ nodes have } -1\}\\
&= (1-\alpha)^{d_\mathsf{c}-1} \tfrac{(1+(1-2x_\ell)^{d_\mathsf{c}-1})}{2}.
\end{aligned}$$

The probability that a check-to-variable message is "?" is the complement of the probability that a check node computes ± 1. Define p_0 to be

$$1-(1-\alpha)^{d_\mathsf{c}-1}. \tag{2}$$

Consider a random variable $V \sim \text{Binomial}(d_\mathsf{v}-1, 1-\alpha)$ with probability mass function $p_V(v)$, capturing the distribution of number of check nodes connected to a variable node. Define p_{+1} and p_{-1} such that

$$p_{+1} = (1-\alpha)^{d_\mathsf{c}-1} \tfrac{(1+(1-2x_\ell)^{d_\mathsf{c}-1})}{2} \tag{3}$$

and

$$p_{-1} = (1-\alpha)^{d_\mathsf{c}-1} \tfrac{(1-(1-2x_\ell)^{d_\mathsf{c}-1})}{2}. \tag{4}$$

Now consider $x_{\ell+1}$, the error probability at a variable node at the $(\ell+1)$th iteration. The fraction of incorrect values held at this variable node is the sum of the probability of two events. The first event is that the message received from the channel is correct, and none of the incoming messages from the connected check nodes is correct, but not all of them are "?", and not only one says different while others are "?". The second event is that the message received from the channel is wrong, and at least one of the incoming messages from the connected check nodes is wrong or at most one check node is correct while all others are "?".

The probability of the first event is:

$$\begin{aligned}
&\mathbb{E}_V\Big[(1-\varepsilon)\big[\Pr\{\text{no connected check nodes sends } 1\}\\
&\quad - \Pr\{\text{all } V \text{ connected check nodes send "?"}\}\\
&\quad - \Pr\{\text{one check node sends } -1 \text{ while others send "?"}\}\big]\Big]\\
&= \sum_{v=1}^{d_\mathsf{v}-1} p_V(v)(1-\varepsilon)[(p_{-1}+p_0)^v - p_0^v - p_{-1}{p_0}^{v-1}].
\end{aligned}$$

The probability of the second event is:

$$\begin{aligned}
&\mathbb{E}_V\Big[\varepsilon\big[\Pr\{\text{at least one connected check nodes send } -1\}\\
&\quad + \Pr\{\text{all } V \text{ connected check nodes send "?"}\}\\
&\quad + \Pr\{\text{one check node sends } +1 \text{ while others send "?"}\}\big]\Big]\\
&= \sum_{v=0}^{d_\mathsf{v}-1} p_V(v)\varepsilon[1-(p_{+1}+p_0)^v + p_0^v + p_{+1}{p_0}^{v-1}].
\end{aligned}$$

Let $x_{\ell+1} = f_{DE}(x_\ell, \varepsilon, \alpha)$, and take the expectation of V according to the binomial distribution to get

$$\begin{aligned}
x_{\ell+1} =& f_{DE}(x_\ell, \varepsilon, \alpha)\\
=& \varepsilon\alpha^{d_\mathsf{v}-1} + \sum_{v=1}^{d_\mathsf{v}-1} \binom{d_\mathsf{v}-1}{v} (1-\alpha)^v \alpha^{(d_\mathsf{v}-1-v)}\\
&\cdot \Big[(1-\varepsilon)[(p_{-1}+p_0)^v - p_0^v - p_{-1}{p_0}^{v-1}]\\
&\quad + \varepsilon[1-(p_{+1}+p_0)^v + p_0^v + p_{+1}{p_0}^{v-1}]\Big].
\end{aligned}$$

To extend to irregular LDPC ensembles, we take the average of the check node distribution and get:

$$p_{+1}^{(irr)} = \rho(1-\alpha)\tfrac{1-\rho(1-2x_\ell)}{2} \text{ and} \tag{5}$$

$$p_{-1}^{(irr)} = \rho(1-\alpha)\tfrac{1+\rho(1-2x_\ell)}{2}. \tag{6}$$

The terms in $f_{DE}(x_\ell, \varepsilon, \alpha)$ have to be averaged over the variable node degree distribution of d_v with function $\lambda(\cdot)$.

B. Fixed Points

It can be seen that $f_{DE}(x, \varepsilon, \alpha)$ is monotonic in x for a set of given α and ε. Hence, by the same arguments as for peeling decoders, for any initial $0 \le \varepsilon = x_0 \le 0.5$, x_ℓ converges to the nearest fixed point of the density evolution equation. We use $\tau_1 \le \tau_2 \le \tau_3 \le \cdots$ to denote these fixed points.

Note that for all $\varepsilon > 0, \alpha > 0$, and $x_\ell = 0$, $f_{DE}(x_\ell, \varepsilon, \alpha) = x_{\ell+1} > 0$. This implies a result similar to Lem. 2 here. With the existence of channel noise and missing wiring, the decoding probability cannot be driven to 0. It is easy to show that for $\varepsilon = 0$, $f_{DE}(x, 0, \alpha)$ has one fixed point at $\tau_1 = 0$. We then focus on the case where $0 < \varepsilon < 0.5$ for the following analysis.

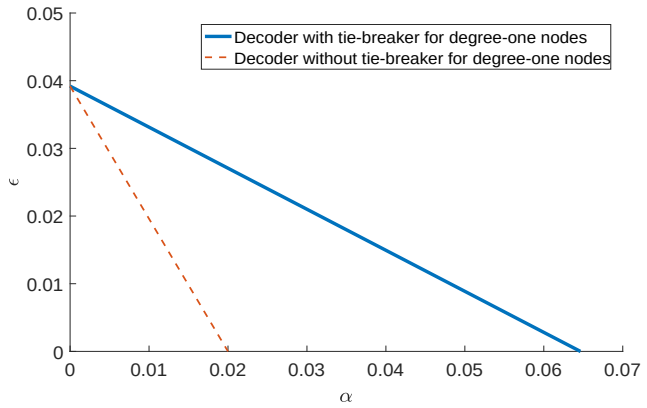

Fig. 4. Decoding a $C^{\infty}(3,6)$ regular LDPC code with α-missing wire Gallager A decoding algorithm over BSC(ε). The useful region where it is beneficial to use decoder is between the curve and α-axis.

Define $p^{+}(x) = (p_{-1}+p_0)^v - p_0^v - p_{-1}{p_0}^{v-1}$ and $p^{-}(x) = 1-(p_{+1}+p_0)^v + p_0^v$. An analytical expression for the channel threshold is the root (τ_2) of the following expression between 0 and 0.5:

$$x\lambda(\alpha) + \lambda\Big(p^{+}(x) - xp^{+}(x) + xp^{-}(x)\Big) = x.$$

The solid line in Fig. 4 shows the useful region of decoding for a $(3,6)$ regular LDPC code with missing connections, which is between τ_1 and τ_2 due to the monotonicity of function f_{DE}. Compared to [13, Fig. 2] where computation at each node is noisy with probability α, the useful region of a decoder with missing connection is larger. In this case, decoders with missing connections outperform those with noisy computation. At any node, if the corresponding incoming message is missing rather than noisy with probability α, the node is more likely to send a correct message than an erroneous one.

C. Performance Analysis

Fig. 5 shows η-thresholds for communication over BSC(ε) with a Gallager A decoder with missing connections. Recall that for a $(3,6)$ regular LDPC code with a fault-free Gallager A decoder, the threshold is roughly 0.039 [34]. Note that P_e can be driven to a fairly small number even with missing wires. Decoding is robust to missing connection defects, though less than the peeling decoder over BEC.

As observed in Fig. 5, a phase transition of the decoding threshold ε with the change of missing connection probability α noticed in the peeling decoder also exists here. In contrast to classic settings, there may be degree-one nodes in decoding graphs due to the random missing connections. Hence, a tie-breaker at a variable node is necessary when the only incoming message from a check node is different from the received message from the channel. Since the channel message is more reliable than internal messages when there are missing connections in the decoder, we choose not to flip the channel message when the only incoming non-erasure message is the opposite. With this minor twist, the decoder useful region increases significantly, as shown in Fig. 4, where the dotted line shows the useful region of the decoder without the tie-breaker for degree-one case, choosing to flip the channel symbol when all incoming non-erasure messages are different from the channel symbol.

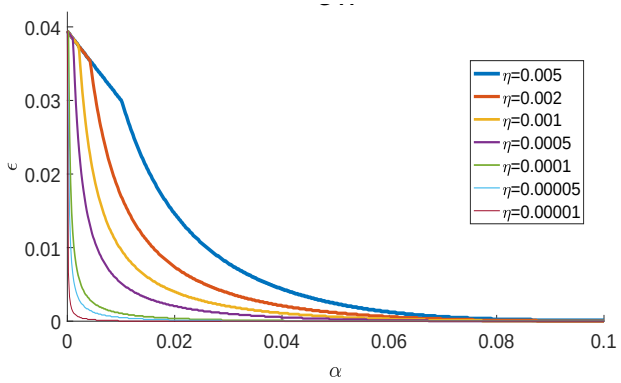

Fig. 5. η-thresholds for decoding a $C^{\infty}(3,6)$ regular LDPC code with α-missing wire Gallager A decoding algorithm over BSC(ε).

VI. Gallager B Decoder

Gallager B decoders are usually more robust than Gallager A decoders without missing connections [18], so we modify the Gallager B algorithm by introducing erasure symbols for missing connections. In the Gallager B decoder, a check node performs the same operation with incoming variable-to-check node messages as Gallager A in Sec. V, sending an unknown symbol "?" if one of the incoming messages is from a disconnected node. At a variable node however, instead of flipping the current value u only when all the incoming messages from connected nodes say $-u$, a variable node in the Gallager B decoder decides to correct the current value u when there are more than b number of incoming messages that are $-u$. This threshold can be iteration-specific to reach optimality. Here, we fix the majority criterion, $b^{*} = \lfloor (d_{\mathsf{v}}+1)/2 \rfloor$, in all iterations because this threshold results in small error probability independent of iteration number in fault-free Gallager B work [35, Sec. 5]. We also choose b^{*} based on the designed code without counting the number of actually connected nodes for simplicity, and it is verified numerically that there is no significant difference in performance.

Similar to the Gallager A model developed in Sec. V, the codeword symmetry conditions C1–C4 are all satisfied in invoking Prop. 1.

A. Density Evolution Equation

The density evolution equation for the Gallager B decoder is similar to Gallager A. Consider a regular $(d_{\mathsf{v}}, d_{\mathsf{c}})$ LDPC code and all-one codeword transmitted over BSC. At iteration ℓ, the probability of a check-to-variable message is "?", $+1$ or -1 with probabilities p_0, p_{+1} and p_{-1}, respectively, which have the same expressions as in Sec. V.

Now consider $x_{\ell+1}$, the error probability at a variable node at the $(\ell+1)$th iteration. The fraction of incorrect values held at this variable node is the sum of the probability of two events. The first event is that the message received from the channel

is correct, and at least $b = \lfloor (d_{\mathrm{v}}+1)/2 \rfloor$ check nodes are connected and send incorrect messages. The second event is that the message received from the channel is wrong, and at most $b-1 = \lfloor (d_{\mathrm{v}}-1)/2 \rfloor$ of the incoming messages from the check nodes are correct. Consider a random variable $V \sim \text{Binomial}(d_{\mathrm{v}}-1, 1-\alpha)$ capturing the distribution of the number of check nodes connected to a variable node.

The probability of the first event is:

$$\begin{aligned}&\mathbb{E}_V\left[(1-\varepsilon)\Pr\{\text{at least } b \text{ check nodes are connected and send} - 1\}\right]\\ &= \sum_{v=b}^{d_{\mathrm{v}}-1} p_V(v)(1-\varepsilon){p_{-1}}^{v}(1-p_{-1})^{d_{\mathrm{v}}-1-v}.\end{aligned}$$

The probability of the second event is:

$$\begin{aligned}&\mathbb{E}_V\left[\varepsilon \Pr\{\text{at most } (b-1) \text{ check nodes send} + 1\}\right]\\ &= \mathbb{E}_V\left[\varepsilon[1-\Pr\{\text{at least } b \text{ check nodes are connected and send} + 1\}]\right]\\ &= \sum_{v=b}^{d_{\mathrm{v}}-1} p_V(v)\varepsilon[1-{p_{+1}}^{v}(1-p_{+1})^{d_{\mathrm{v}}-1-v}].\end{aligned}$$

Taking the expectation of V according to the binomial distribution, we have

$$\begin{aligned}x_{\ell+1} = \sum_{v=b}^{d_{\mathrm{v}}-1}\binom{d_{\mathrm{v}}-1}{v}(1-\alpha)^{v}\alpha^{(d_{\mathrm{v}}-1-v)}\Big[&(1-\varepsilon)[{p_{-1}}^{v}\\ &\cdot(1-p_{-1})^{d_{\mathrm{v}}-1-v}] + \varepsilon[1-{p_{+1}}^{v}(1-p_{+1})^{d_{\mathrm{v}}-1-v}]\Big].\end{aligned}$$

The density evolution equation can also be extended to irregular LDPC codes, with changes in parameters $b(x) = \lfloor (\lambda(x)+1)/2 \rfloor$, $p_{+1}^{(irr)}$, and $p_{-1}^{(irr)}$ defined in expressions (5) and (6).

B. Performance Analysis

We carry out detailed performance characterization of a Gallager B decoder with missing connections and show that such a decoder is indeed more robust to missing connections than Gallager A.

Note when variable node degree $d_{\mathrm{v}} = 3$ for a regular LDPC code, a fault-free Gallager B decoder with the defined threshold $b = \lfloor (d_{\mathrm{v}}+1)/2 \rfloor$ is equivalent to a fault-free Gallager A decoder. However, due to the modification the of Gallager A decoder to keep the received channel message when there exists only one incoming message, in the case of missing connections, these two decoders behave differently for decoding a $C^{\infty}(3,6)$ regular LDPC code.

One interesting phenomenon shown in Fig. 6 is that decoding thresholds first increase with increasing decoder missing connection probability. This error enhancement phenomenon is introduced by the missing connections, essentially resulting in a change of choice for threshold b in each iteration to achieve a lower error rate. This SF phenomenon demonstrates that optimization of degree distribution and threshold b in each iteration can be utilized to combat missing connections. A similar SF result shows that the errors introduced in estimating

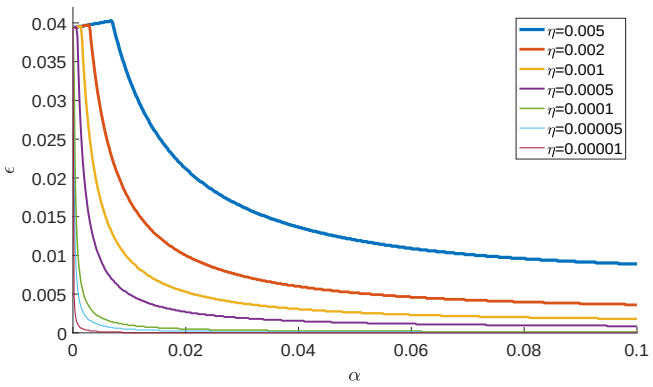

Fig. 6. η-thresholds for decoding a $C^{\infty}(3,6)$ regular LDPC code with α-missing wire Gallager B decoding algorithm over BSC(ε).

Markov random field models can be partially canceled and benefit end-to-end inference performance [36]. SF effects due to noise in computational elements, rather than graphical model structure errors as here, have been observed in [27], [29], [37] and later specifically in LDPC decoders [23], [25].

VII. Moving towards Practice

Though performance analysis of LDPC decoders with missing connections using density evolution is an important topic in coding theory, our eventual goal is to use analytical understanding for practical system design.

Towards this end, we first briefly discuss how one can use DE analysis to study sensitivity of codes and decoders, so as to give insight into resource allocation over the entire telecommunications system. In particular we ask whether more resources should be spent in manufacturing or in operation. Second, as DE analysis is an asymptotic approximation of practical finite-length codes, we also perform simulations to understand how well the asymptotics describe finite-length code performance. Finally we note that increasing the accuracy of semiconductor fabrication by just a small amount requires a significant increase in manufacturing cost (which already takes tens of billions of dollars to build facilities, and limits growth of the industry). As such, we perform preliminary manufacturing yield analysis to show potential industrial impact. For brevity, this section is largely restricted to Gallager A.

A. Decoder Sensitivity

Should the industry invest more resources in operating good communication channels or in manufacturing better receiver hardware?

Taking derivatives of the density evolution function $x_{\ell+1} = f_{DE}(x_l, \varepsilon, \alpha)$ with respect to ε and α and evaluating at $\bar{P}_e^{(\infty)} = x_\ell = x_{\ell+1}$, we find the impact of channel noise level and missing connection probability on the final error rate.

$$\begin{aligned}\bar{P}_e^{(\infty)}(\varepsilon,\alpha) = \varepsilon\lambda(\alpha) + \lambda\Big((1-\varepsilon)p^{+}(\bar{P}_e^{(\infty)}(\varepsilon,\alpha))\\ +\varepsilon p^{-}(\bar{P}_e^{(\infty)}(\varepsilon,\alpha))\Big).\end{aligned}$$

Denote

$$g(x) = (1-\varepsilon)p^{+}(x(\varepsilon,\alpha)) + \varepsilon p^{-}(x(\varepsilon,\alpha)).$$

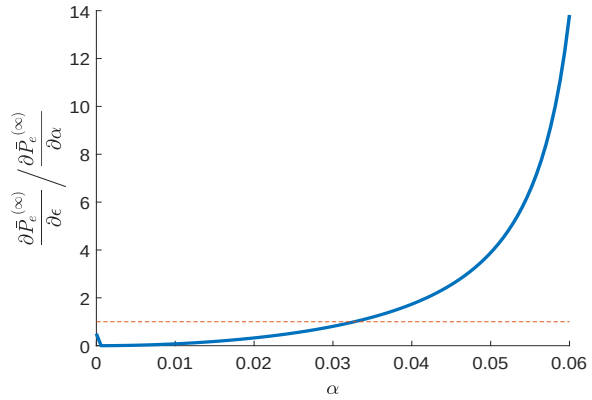

Fig. 7. Comparison between the derivative of $\bar{P}_e{}^{(\infty)}(\varepsilon, \alpha)$ with respect to ε and α of decoding a $C^{\infty}(3, 6)$ regular LDPC code with α-missing wire Gallager A decoding algorithm over BSC(ε), when ε and α are at the boundary of decoder useful region.

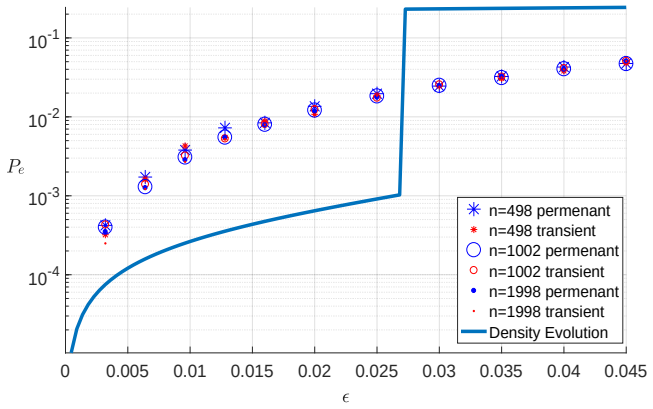

Fig. 8. P_e of decoding $(3, 6)$ LDPC code with finite block-length under Gallager A decoding algorithm over BSC with permanently and transiently missing connection probability $\alpha = 0.02$.

Take partial derivatives of each side with respect to ε:

$$\frac{\partial x(\varepsilon,\alpha)}{\partial \varepsilon} = \lambda(\alpha) + \frac{\partial \lambda(g(x(\varepsilon,\alpha)))}{\partial x(\varepsilon,\alpha)} \frac{\partial x(\varepsilon,\alpha)}{\partial \varepsilon} + \frac{\partial \lambda(g(x(\varepsilon,\alpha)))}{\partial \varepsilon}$$
$$= \frac{\lambda(\alpha) + \frac{\partial \lambda(g(x(\varepsilon,\alpha)))}{\partial \varepsilon}}{1 - \frac{\partial \lambda(g(x(\varepsilon,\alpha)))}{\partial x(\varepsilon,\alpha)}}.$$

Similarly,

$$\frac{\partial x(\varepsilon,\alpha)}{\partial \alpha} = \varepsilon \frac{\partial \lambda(\alpha)}{\partial \alpha} + \frac{\partial \lambda(g(x(\varepsilon,\alpha)))}{\partial x(\varepsilon,\alpha)} \frac{\partial x(\varepsilon,\alpha)}{\partial \alpha} + \frac{\partial \lambda(g(x(\varepsilon,\alpha)))}{\partial \alpha}$$
$$= \frac{\varepsilon \frac{\partial \lambda(\alpha)}{\partial \alpha} + \frac{\partial \lambda(g(x(\varepsilon,\alpha)))}{\partial \alpha}}{1 - \frac{\partial \lambda(g(x(\varepsilon,\alpha)))}{\partial x(\varepsilon,\alpha)}}.$$

Fig. 7 illustrates the ratio of the derivative of $\bar{P}_e{}^{(\infty)}(\varepsilon, \alpha)$ with respect to ε and α, when α and ε are at the boundary of useful region depicted in Fig. 5:

$$\frac{\partial \bar{P}_e{}^{(\infty)}(\varepsilon, \alpha)}{\partial \varepsilon} \Big/ \frac{\partial \bar{P}_e{}^{(\infty)}(\varepsilon, \alpha)}{\partial \alpha}.$$

Different from our intuition, both derivate values are negative at the boundary of the useful region. Recall the linear relationship of ε and α at the boundary of the useful region; with the increase of α, ε has to decrease in order to stay in the useful region, resulting in the decrease in $\bar{P}_e{}^{(\infty)}(\varepsilon, \alpha)$.

When operating at the edge of the useful region, as we can see in Fig. 7, it is advantageous to put resources into circuit manufacturing up to an α value of roughly 0.03 where the curve crosses the equal-ratio point, whereas it is advantageous to put resources into the channel thereafter. Thus aiming for manufacturing that achieves such a crossover point α may be an appropriate resource allocation strategy.

B. Finite-Length Simulation

We simulate finite-length systems having decoders with either transiently or permanently missing connections, to demonstrate performance is comparable in the two settings and predicted by density evolution. For $(3, 6)$-regular LDPC codes with blocklength $n = 498$, 1002, and 1998 drawn at random from the code ensemble using socket-switching, we randomly simulate decoding performance with connections either permanently removed before the decoding starts or transiently removed during each decoding iteration for various sets of (ε, α). For each trial, decoding is performed for more than 30 iterations (error probability usually convergences within 10 iterations in fault-free decoding). For each channel noise level and missing connection probability, the decoding error probability is averaged over 100 randomly selected code realizations and missing connection realizations.

As Fig. 8 illustrates, the performance of finite-length codes resembles the asymptotic performance of codes. As expected, below the decoding threshold, the symbol error rates of finite-length codes are higher than asymptotic performance. For fault-free decoders with channel noise below threshold, the asymptotic symbol error rate is 0, whereas in the case of finite-length codes P_e increases smoothly with increasing ε [38]. Similar to fault-free decoders for finite-length codes, decoders with missing connections show a similar trend of increasing P_e. Further, notice that the performances of transiently and permanently missing connection cases are close to one another. We chose $\alpha = 0.02$ because it is within the defective interconnect range of 1–15% [7]; see also Sec. VII-A. Different from [38], we do not expurgate codes with small stopping sets. Also recall from Thm. 1 that the concentration of the individual performance around the ensemble average is exponential in blocklength and the concentration happens more slowly in the case of missing connections compared to fault-free decoders. Hence, there is more numerical variation in the simulation results at all blocklengths, especially for small n. Nevertheless, simulations show that the asymptotic analysis of decoders with missing connections has practical significance.

C. Semiconductor Manufacturing Yield Analysis

By understanding the computational purpose of circuits (here decoding) it is often possible to raise effective manufacturing yield above the raw yield where all components must be fault-free [39], [40]. To demonstrate that the effective yield of LDPC decoder circuits increases by allowing missing connection probability α that still guarantees decoding performance η, we apply the error-tolerant methodology [39],

[40]. Threshold testing in [39] accepts a chip when the chip's performance passes a specified threshold; the increase in effective yield is the amount of chips with defects but still meeting the performance requirement. For decoders with missing connections, this threshold is the maximum symbol error rate η. We want to find the highest missing connection probability $\alpha_{max}(G^n, \varepsilon, \eta)$ such that for every decoder with $\alpha \leq \alpha_{max}$, the resulting P_e is under the target error rate η for a given code ensemble G^n and channel quality ε.

Let $\phi(\alpha)$ be the yield factor, the expected percentage of decoders with missing connection probability α, and $p(\alpha)$ be the probability that the circuit has defect density α, often taken as an exponential distribution [40]. Then the effective yield is:

$$Y = \int_0^{\alpha_{max}(G^n,\varepsilon,\eta)} p(\alpha)\phi(\alpha)d\alpha.$$

For the $\mathcal{C}^\infty(3,6)$ LDPC code ensemble and $\eta = 10^{-5}$, for a large range of possible channel values ε, Fig. 3 shows us that $\alpha_{max} = 0.01$ is more than sufficient for the case of peeling decoder under BEC. It is straightforward to see that, compared to the yield of the fault-free case $Y_0 = p(0)\phi(0)$, allowing some error-tolerance in manufacturing may increase effective yield significantly. For the exponential distribution function, the absolute increase in yield is linear in α_{max} [40]:

$$\Delta Y = Y - Y_0 = \alpha_{max}\tfrac{D_0 A}{(1+AD_0)^2},$$

where D_0 is the defect density (average number of defects per unit of chip area), and A is the chip area. Likewise the fractional increase in yield, is:

$$\Delta Y/Y_0 = \tfrac{Y-Y_0}{Y_0} = \alpha_{max}\tfrac{D_0 A}{1+AD_0}.$$

As shown in the previous sections, a small defect rate α does not degrade the performance too much. However, as reported in the semiconductor manufacturing industry, a 1% reduction in yield can result in a 12% reduction in profit [41], [42]. Hence even allowing a small probability of defects α can save a significant amount of wastage and cost without much change in performance.

VIII. Conclusion

This paper investigated the performance of message-passing decoders with transiently and permanently missing connections that might be caused by process variation in manufacturing or timing errors in intra-chip communications (or both). We derived density evolution equations to characterize the error probability in the peeling decoder over the BEC and modifications of the Gallager A and Gallager B decoders over the BSC, using erasure symbols to represent missing connections. Although the error probability cannot be driven to 0 in the presence of missing connections, it can be suppressed to a small value η when the channel noise level is under a certain decoding threshold ε^*. That is, η-reliable communication is possible with faulty decoders with missing connections. In a sense, even when the encoder and decoder speak different languages, the result is not catastrophic. A novel structural stochastic facilitation is also observed in Gallager B decoders with missing connections.

Future work involves considering not just decoders with missing connections, but also miswired and noisy decoders. One may also design new decoder architectures to ensure reliable communication even with miswiring; for example, horizontal connections, a crucial structure in the cortex contributing to the filling in of missing parts in visual images [43, Ch. 8.33], can be added to decoder designs. Code optimization and new decoding algorithms can also be utilized to take advantage of the stochastic facilitation effect.

Acknowledgment

The authors thank A. Patil and N. R. Shanbhag for discussions on nanoscale circuits, anonymous reviewers for helpful comments, and A. C. Singer and J. Cowan for encouragement.

Appendix A
Probability Theory Definitions

Before diving into the proof of Thm. 1, some probability theory definitions and the Hoeffding-Azuma inequality are reviewed here. Consider a space $(\Omega, \mathcal{F})$, where Ω is a sample space, and a σ-algebra $\mathcal{F}$ contains subsets of Ω. A random variable Z is an $\mathcal{F}$-measurable function from a probability space into the real number. If there is a collection $(Z_\gamma|\gamma \in C)$ of random variables $Z_\gamma : \Omega \to \mathbb{R}$, then

$$\mathcal{Z} = \sigma(Z_\gamma|\gamma \in C)$$

is defined to be the smallest σ-algebra $\mathcal{Z}$ on Ω such that each map $(Z_\gamma|\gamma \in C)$ is $\mathcal{Z}$-measurable.

Definition 1 (Filtration): Let $\{\mathcal{F}_i\}$ be a sequence of σ-algebras with respect to the same sample space Ω. These $\mathcal{F}_i$ are said to form a *filtration* if $\mathcal{F}_0 \subseteq \mathcal{F}_1 \subseteq \cdots$ are ordered by refinement in the sense that each subset of Ω in $\mathcal{F}_i$ is also in $\mathcal{F}_j$ for $i \leq j$. Also $\mathcal{F}_0 = \{\emptyset, \Omega\}$.

The conditional expectation of a random variable Z given a σ-algebra $\mathcal{F}$ is a random variable denoted by $E[Z|\mathcal{F}]$.

Definition 2 (Martingale): Let $\mathcal{F}_0 \subseteq \mathcal{F}_1 \subseteq \cdots$ be a filtration on Ω and let $Z_0, Z_1, \ldots$ be a sequence of random variables on Ω such that Z_i is $\mathcal{F}_i$-measurable. Then $Z_0, Z_1, \ldots$ is a *Martingale* with respect to the filtration $\mathcal{F}_0 \subseteq \mathcal{F}_1 \subseteq \cdots$ if $E[Z_i|\mathcal{F}_{i-1}] = Z_{i-1}$.

Definition 3 (Doob's Martingale): Let $\mathcal{F}_0 \subseteq \mathcal{F}_1 \subseteq \cdots$ be a filtration on Ω and let Z be a random variable on Ω. Then the sequence of random variables $Z_0, Z_1, \ldots$ such that $Z_i = E[Z|\mathcal{F}_i]$ is a Doob's Martingale.

Lemma 3 (Hoeffding-Azuma Inequality [31], [44], [45]): Let $Z_0, Z_1, \ldots$ be a Martingale with respect to the filtration $\mathcal{F}_0 \subseteq \mathcal{F}_1 \subseteq \cdots$ such that for each $i > 0$, the following bounded difference condition is satisfied

$$|Z_i - Z_{i-1}| \leq \alpha_i,\ \alpha_i \in [0, \infty).$$

Then for all $n > 0$ and any $\xi > 0$,

$$\Pr\left[|Z_n - Z_0| \geq \xi\right] \leq 2\exp\left(-\frac{\xi^2}{2\sum_{k=1}^n \alpha_k^2}\right).$$

APPENDIX B
CONCENTRATION: PERMANENTLY MISSING CONNECTIONS

The proof of Thm. 1 is an extension from and largely identical to [13, Thm. 2], [31, Thm. 2], or [32, Thm. 4.94]. We want to construct a Doob's Martingale with respect to the fraction of error held on each edge during the random revealing process and to show that the difference of the object of interest between each iteration is bounded by a number not related to the number of iterations.

Recall Z denotes the number of incorrect values held at the end of the ℓth iteration for a specific $(g, y, w) \in \Omega$, where g is a specific bipartite Tanner graph to represent the choice of LDPC code with variable node degree $d_{\mathtt{v}}$ and check node degree $d_{\mathtt{c}}$, y is a specific input to the decoder, w is a particular realization of the decoder with missing wires, and Ω is the sample space. Let $\equiv_i$, $0 \leq i \leq m$ be a sequence of equivalence relations on Ω ordered by refinement, such that $(g', y', w') \equiv_i (g'', y'', w'')$ implies $(g', y', w') \equiv_{i-1} (g'', y'', w'')$. The equivalence relations define equivalence classes by partial equalities such that $(g', y', w', u') \equiv_i (g'', y'', w'', u'')$ *if and only if* the realizations of random quantities revealed in the first i steps for both pairs is the same.

Next we use the technique of exposing the edges in the decoding graph in sequence. The first case is when wires are permanently missing. Note that even with positive probability of missing connections α, for a specific code realization, the number of potentially connected edges can be at most $nd_{\mathtt{v}}$. Hence, we expose at most $nd_{\mathtt{v}}$ edges one at a time. At step $i \leq nd_{\mathtt{v}}$, we expose the particular check node socket that is connected to the ith variable node socket. Next, in the following n steps, we expose the received values y_i from the channel one at a time. At the end of the $n(d_{\mathtt{v}} + 1)$ steps, the decoder missing wire probability is also realized, since the defect is permanent. Then we have $(g', y', w') \equiv_i (g'', y'', w'')$ if and only if the information revealed in the first i steps for both pairs is the same.

Now, define $Z_0, Z_1, ..., Z_m$ by

$$Z_i(g, y, w) = E[Z(g', y', w')|(g', y', w') \equiv_i (g, y, w)],$$

where $Z_0 = E[Z]$ and $Z_m = Z$. By construction, $Z_0, Z_1, ..., Z_m$ is a Doob's Martingale. We then use Lem. 3 to give bounds on

$$\Pr[|Z - E[Z]| > nd_{\mathtt{v}}\varepsilon/2] = \Pr[|Z_m - Z_0| > nd_{\mathtt{v}}\varepsilon/2].$$

To use Azuma's inequality, we first need to prove that for each consecutive member in the sequence $Z_0, Z_1, ..., Z_m$, the difference is bounded:

$$|Z_{i+1}(g, y, w) - Z_i(g, y, w)| \leq \delta_i, i = 0, 1, ..., m - 1$$

where δ_i depends on $d_{\mathtt{v}}, d_{\mathtt{c}}$, and ℓ.

It was shown in [31] that for the fault-free decoder without any missing wire, when edges are exposed,

$$|Z_{i+1}(g, y, w) - Z_i(g, y, w)| \leq 8(d_{\mathtt{v}}d_{\mathtt{c}})^{\ell}, 0 \leq i \leq nd_{\mathtt{v}}.$$

In our case when there exist permanently missing connections, the difference when exposing edges is that the number of edges existing is smaller, and bounded by $nd_{\mathtt{v}}$. The expected number of edges left is $nd_{\mathtt{v}}(1 - \alpha)$. The bound established above still holds with a change of the steps number:

$$|Z_{i+1}(g, y, w) - Z_i(g, y, w)| \leq 8(d_{\mathtt{v}}d_{\mathtt{c}})^{\ell}, 0 \leq i \leq nd_{\mathtt{v}}.$$

It was also shown that when channel outputs are revealed, the difference in each element in the sequence is bounded by

$$|Z_{i+1}(g, y, w) - Z_i(g, y, w)| \leq 2(d_{\mathtt{v}}d_{\mathtt{c}})^{\ell},$$

where $nd_{\mathtt{v}} \leq i \leq n(d_{\mathtt{v}} + 1)$ in the case where some wires are permanently missing. Then the theorem follows from applying Azuma's inequality to the Martingale constructed.

APPENDIX C
CONCENTRATION: TRANSIENTLY MISSING CONNECTIONS

The second case is when wires are transiently missing at each decoding iteration. The Martingale is constructed differently. Instead of exposing edges, at ℓ iterations, we sequentially expose the realization of edges at different iterations. Since each edge is missing independently from others with probability α, only sockets whose nodes are connected through these edges are affected. In each iteration, there are two realizations for each edge (present or missing), then for all previous ℓ iterations, the total number affected edges is bounded by $2(2d_{\mathtt{v}}d_{\mathtt{c}})^{\ell}$. With symmetry of switching node sockets:

$$|Z_{i+1}(g, y, w) - Z_i(g, y, w)| \leq 8(2d_{\mathtt{v}}d_{\mathtt{c}})^{\ell}$$

where $n(d_{\mathtt{v}} + 1) \leq i \leq m$.

Hence, in the transiently missing wire case, the bounded difference $\delta_i = 8(2d_{\mathtt{v}}d_{\mathtt{c}})^{\ell}$. The theorem follows from applying Azuma's inequality to the Martingale constructed.

REFERENCES

[1] L. Chang, A. Chatterjee, and L. R. Varshney, "LDPC decoders with missing connections," in *Proc. 2016 IEEE Int. Symp. Inf. Theory*, Jul. 2016, pp. 1576–1580.

[2] R. G. Gallager, *Low-Density Parity-Check Codes*. Cambridge, MA: MIT Press, 1963.

[3] A. A. Al-Yamani, S. Ramsundar, and D. K. Pradhan, "A defect tolerance scheme for nanotechnology circuits," *IEEE Trans. Circuits Syst. I*, vol. 54, no. 11, pp. 2402–2409, Nov. 2007.

[4] S.-L. Jeng, J.-C. Lu, and K. Wang, "A review of reliability research on nanotechnology," *IEEE Trans. Rel.*, vol. 56, no. 3, pp. 401–410, Sep. 2007.

[5] Y. M. Chee and A. C. H. Ling, "Limit on the addressability of fault-tolerant nanowire decoders," *IEEE Trans. Comput.*, vol. 58, no. 1, pp. 60–68, Jan. 2009.

[6] J. R. Heath, P. J. Kuekes, G. S. Snider, and R. S. Williams, "A defect-tolerant computer architecture: Opportunities for nanotechnology," *Science*, vol. 280, no. 5370, pp. 1716–1721, Jun. 1998.

[7] M. Haselman and S. Hauck, "The future of integrated circuits: A survey of nanoelectronics," *Proc. IEEE*, vol. 98, no. 1, pp. 11–38, Jan. 2010.

[8] M. A. Breuer, S. K. Gupta, and T. M. Mak, "Defect and error tolerance in the presence of massive numbers of defects," *IEEE Des. Test. Comput.*, vol. 21, no. 3, pp. 216–227, May-June 2004.

[9] S. Ghosh and K. Roy, "Parameter variation tolerance and error resiliency: New design paradigm for the nanoscale era," *Proc. IEEE*, vol. 98, no. 10, pp. 1718–1751, Oct. 2010.

[10] M. M. Mansour and N. R. Shanbhag, "A 640-Mb/s 2048-bit programmable LDPC decoder chip," *IEEE J. Solid-State Circuits*, vol. 41, no. 3, pp. 684–698, Mar. 2006.

[11] P. Elias, "Computation in the presence of noise," *IBM J. Res. Develop.*, vol. 2, no. 4, pp. 346–353, Oct. 1958.

[12] S. Winograd and J. D. Cowan, *Reliable Computation in the Presence of Noise*. Cambridge, MA: MIT Press, 1963.

[13] L. R. Varshney, "Performance of LDPC codes under faulty iterative decoding," *IEEE Trans. Inf. Theory*, vol. 57, no. 7, pp. 4427–4444, Jul. 2011.

[14] S. M. S. Tabatabaei Yazdi, H. Cho, and L. Dolecek, "Gallager B decoder on noisy hardware," *IEEE Trans. Commun.*, vol. 61, no. 5, pp. 1660–1673, May 2013.

[15] S. M. S. Tabatabaei Yazdi, C.-H. Huang, and L. Dolecek, "Optimal design of a Gallager B noisy decoder for irregular LDPC codes," *IEEE Commun. Lett.*, vol. 16, no. 12, pp. 2052–2055, Dec. 2012.

[16] F. Leduc-Primeau and W. J. Gross, "Faulty Gallager-B decoding with optimal message repetition," in *Proc. 50th Annu. Allerton Conf. Commun. Control Comput.*, Oct. 2012, pp. 549–556.

[17] F. Leduc-Primeau, F. R. Kschischang, and W. J. Gross, "Energy optimization of LDPC decoder circuits with timing violations," in *Proc. IEEE Int. Conf. Commun. (ICC 2015)*, Jun. 2015, pp. 412–417.

[18] C.-H. Huang, Y. Li, and L. Dolecek, "Gallager B LDPC decoder with transient and permanent errors," *IEEE Trans. Commun.*, vol. 62, no. 1, pp. 15–28, Jan. 2014.

[19] C. Kameni Ngassa, V. Savin, E. Dupraz, and D. Declercq, "Density evolution and functional threshold for the noisy min-sum decoder," *IEEE Trans. Commun.*, vol. 63, no. 5, pp. 1497–1509, May 2015.

[20] B. Vasić and S. K. Chilappagari, "An information theoretical framework for analysis and design of nanoscale fault-tolerant memories based on low-density parity-check codes," *IEEE Trans. Circuits Syst. I*, vol. 54, no. 11, pp. 2438–2446, Nov. 2007.

[21] E. Dupraz, D. Declercq, B. Vasić, and V. Savin, "Analysis and design of finite alphabet iterative decoders robust to faulty hardware," *IEEE Trans. Commun.*, vol. 63, no. 8, pp. 2797–2809, Aug. 2015.

[22] O. A. Rasheed, P. Ivaniš, and B. Vasić, "Fault-tolerant probabilistic gradient-descent bit flipping decoder," *IEEE Commun. Lett.*, vol. 18, no. 9, pp. 1487–1490, Sep. 2014.

[23] B. Vasić, P. Ivaniš, S. Brkic, and V. Ravanmehr, "Fault-resilient decoders and memories made of unreliable components," in *Proc. 2015 Inf. Theory Appl. Workshop*, Feb. 2015.

[24] L. R. Varshney, "Toward limits of constructing reliable memories from unreliable components," in *Proc. IEEE Inf. Theory Workshop (ITW'15)*, Oct. 2015, pp. 114–118.

[25] S. Brkic, P. Ivaniš, and B. Vasić, "Guaranteed error correction of faulty bit-flipping decoders under data-dependent gate failures," in *Proc. 2016 IEEE Int. Symp. Inf. Theory*, Jul. 2016, pp. 1561–1565.

[26] C.-H. Huang, Y. Li, and L. Dolecek, "Belief propagation algorithms on noisy hardware," *IEEE Trans. Commun.*, vol. 63, no. 1, pp. 11–24, Jan. 2015.

[27] A. Karbasi, A. H. Salavati, A. Shokrollahi, and L. R. Varshney, "Noise facilitation in associative memories of exponential capacity," *Neural Comput.*, vol. 26, no. 11, pp. 2493–2526, Nov. 2014.

[28] S. Brkic, O. A. Rasheed, P. Ivaniš, and B. Vasić, "On fault-tolerance of the Gallager B decoder under data-dependent gate failures," *IEEE Commun. Lett.*, vol. 19, no. 8, pp. 1299–1302, Aug. 2015.

[29] H. Chen, L. R. Varshney, and P. K. Varshney, "Noise-enhanced information systems," *Proc. IEEE*, vol. 102, no. 10, pp. 1607–1621, Oct. 2014.

[30] P. Ivaniš, B. Vasić, and D. Declercq, "Performance evaluation of faulty iterative decoders using absorbing Markov chains," in *Proc. 2016 IEEE Int. Symp. Inf. Theory*, Jul. 2016, pp. 1566–1570.

[31] T. J. Richardson and R. L. Urbanke, "The capacity of low-density parity-check codes under message-passing decoding," *IEEE Trans. Inf. Theory*, vol. 47, no. 2, pp. 599–618, Feb. 2001.

[32] T. Richardson and R. L. Urbanke, *Modern Coding Theory*. Cambridge: Cambridge University Press, 2008.

[33] H. Mori and T. Wadayama, "Performance analysis based on density evolution on fault erasure belief propagation decoder," in *Proc. 2016 IEEE Int. Symp. Inf. Theory*, Jul. 2016, pp. 1571–1575.

[34] L. Bazzi, T. J. Richardson, and R. L. Urbanke, "Exact thresholds and optimal codes for the binary-symmetric channel and Gallager's decoding algorithm A," *IEEE Trans. Inf. Theory*, vol. 50, no. 9, pp. 2010–2021, Sep. 2004.

[35] V. Guruswami, "Iterative decoding of low-density parity check codes (a survey)," arXiv:cs/0610022 [cs.IT]., 2006.

[36] M. J. Wainwright, "Estimating the "wrong" graphical model: Benefits in the computation-limited setting," *J. Mach. Learn. Res.*, vol. 7, pp. 1829–1859, Dec. 2006.

[37] T. J. Hamilton, S. Afshar, A. van Schaik, and J. Tapson, "Stochastic electronics: A neuro-inspired design paradigm for integrated circuits," *Proc. IEEE*, vol. 102, no. 5, pp. 843–859, May 2014.

[38] A. Amraoui, A. Montanari, T. Richardson, and R. L. Urbanke, "Finite-length scaling for iteratively decoded LDPC ensembles," *IEEE Trans. Inf. Theory*, vol. 55, no. 2, pp. 473–498, Feb. 2009.

[39] M. A. Breuer and H. Zhu, "An illustrated methodology for analysis of error tolerance," *IEEE Des. Test. Comput.*, vol. 25, no. 2, pp. 168–177, Mar.-Apr. 2008.

[40] Z. Jiang and S. K. Gupta, "Threshold testing: Improving yield for nanoscale VLSI," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 28, no. 12, pp. 1883–1895, Dec. 2009.

[41] K. Flamm, "The impact of DRAM design innovation on manufacturing profitability," *Future Fab Int.*, no. 35, Nov. 2010.

[42] S. Mittal, "A survey of architectural techniques for managing process variation," *ACM Comput. Surv.*, vol. 48, no. 54, May 2016.

[43] P. Brodal, *The Central Nervous System*. New York, NY: Oxford University Press, 2016.

[44] K. Azuma, "Weighted sums of certain dependent random variables," *Tohoku Math. J.*, vol. 19, no. 3, pp. 357–367, 1967.

[45] W. Hoeffding, "Probability inequalities for sums of bounded random variables," *J. Am. Stat. Assoc.*, vol. 58, no. 301, pp. 13–30, Mar. 1963.

Linjia Chang (S'14) received the B.S. degree with honors and the M.S. degree in electrical and computer engineering from the University of Illinois at Urbana-Champaign in 2014 and 2016 respectively. Her research interests span stochastic information processing systems, coding theory, networks, and data analytics. She is a member of Eta Kappa Nu.

Avhishek Chatterjee received the Ph.D. degree in electrical and computer engineering from the University of Texas at Austin in 2015. He is currently a post doctoral research associate with the Coordinate Science Laboratory, University of Illinois at Urbana-Champaign. His research interests lie in theoretical studies of dynamics, optimal designs, and operations of stochastic networks. He is involved in dynamics and inference in social networks, fundamental limits and optimal operations of crowdsourcing systems, resource allocation and dynamics in communication and computer networks, and fundamental limits and resource allocation in nanoscale circuits.

Lav R. Varshney (S'00–M'10–SM'15) received the B.S. degree (*magna cum laude*) in electrical and computer engineering with honors from Cornell University, Ithaca, New York, in 2004. He received the S.M., E.E., and Ph.D. degrees, all in electrical engineering and computer science, from the Massachusetts Institute of Technology, Cambridge, in 2006, 2008, and 2010, where his theses received the E. A. Guillemin Thesis Award and the J.-A. Kong Award Honorable Mention.

He is an assistant professor in the Department of Electrical and Computer Engineering, the Coordinated Science Laboratory, the Beckman Institute, and the Neuroscience Program at the University of Illinois at Urbana-Champaign. During 2010–2013, he was a research staff member at the IBM Thomas J. Watson Research Center, Yorktown Heights, New York. His research interests include information and coding theory; limits of nanoscale, human, and neural computing; human decision making and collective intelligence; and creativity.

Dr. Varshney is a member of Eta Kappa Nu, Tau Beta Pi, and Sigma Xi. He received the IBM Faculty Award in 2014 and was a Finalist for the Bell Labs Prize, in 2014 and 2016. He and his students have won several best paper awards. His work appears in the anthology, *The Best Writing on Mathematics 2014* (Princeton University Press). He is a founding member of the IEEE Special Interest Group on Big Data in Signal Processing and currently serves on the Shannon Centenary Committee of the IEEE Information Theory Society. He also currently serves on the advisory board of the AI XPRIZE.